

allinea



Leaders in parallel software development tools

Large Scale Debugging on Titan and Mira with Allinea DDT

David Lecomber
Allinea Software
david@allinea.com

Why?

Allinea's Tools



Reduce
Development Time





Increase Application
Performance



Understand Application
Behaviour

HPC could be brain surgery

- Brain aneurysms
 - 2-5% of population – most are undiagnosed
 - 30,000 rupture in US each year – 40% fatal
 - Early discovery and treatment increases survival rates
- Neurosurgery as HPC
 - MRI provides the blood vessel structure
 - Intra-cranial blood flow and pressures is just complex CFD
 - Full brain 3D model is 2-10GB geometry



Impact of Petascale and Beyond

- Individualized HPC
 - Patient's MRI scan enables surgical decision: whether to operate, how to operate, ...
 - Circle of Willis requires super-Petascale ~~machine~~ software
 - Need answer in minutes or hours
- Machines can do 20 PetaFLOPs
 - Super-Petascale will be affordable soon
 - Software has to scale

Real scaling challenge

- Crashes at 49,152 cores on Cray XC30
 - Error message “Terminated”. Thanks.
- Now what?
 - Try other (inferior?) partitioner?
 - Invest weeks in bug fix by trial and error?
 - Write own partitioning library?
- Why use a debugger?
 - It’s about time

Allinea DDT 4.2.1-36484

File View Control Search Tools Window Help

Current Group: All Focus on current: ☒ Group ☐ Process ☐ Thread ☐ Step Threads Together

24576 processes (0-24575) Paused: 17223 Playing: 7353 Finished: 0

Currently selected: 260 (on nid00194, pid 9481, main thread IWP 9481)

Create Group

Project Files

Search (Ctrl+K)

VolumeTrav...
wave.c
weird.c
WholeGeom...
Writer.cc
wspace.c
XdrFileWrite...
XdrMemRea...
XdrMemWrit...
XdrReader.c...
XdrWriter.cc
XmlAbstract...
xyzpart.c
External Code

MpiEnvironment.cc xyzpart.c

```
551 ikvsortii(ntsamples, allpicks);
552
553
554 /* Select the final splitters. Set the boundaries to s
555 for (i=1; i<nps; i++)
556     mypicks[i] = allpicks[i*ntsamples/nps];
557 mypicks[0].key = IDX_MIN;
558 mypicks[nps].key = IDX_MAX;
559
560
561 WCOREPOP; /* free allpicks */
562
563 STOPTIMER(ctrl, ctrl->AuxTmr2);
564 STARTTIMER(ctrl, ctrl->AuxTmr3);
565
```

Locals Current Line(s) Current Stack

Current Line(s)

Variable Name	Value
allpicks	0x2aab8055e010
i	2245
my picks	0x2a6f8f0
nps	24575
ntsamples	1818550

Type: none selected

Input/Output Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output Logbook

Stacks

Processes	Threads	Function
17223	17223	main (main.cc:37)
17223	17223	SimulationMaster::SimulationMaster (SimulationMaster.cc:63)
17223	17223	SimulationMaster::Initialise (SimulationMaster.cc:154)
17223	17223	hemelb::geometry::GeometryReader::LoadAndDecompose (GeometryReader.cc:188)
17223	17223	hemelb::geometry::GeometryReader::OptimiseDomainDecomposition (GeometryReader.c
17223	17223	hemelb::geometry::decomposition::OptimisedDecomposition::OptimisedDecomposition
17223	17223	hemelb::geometry::decomposition::OptimisedDecomposition::CallParmetis (Optimised
17223	17223	ParMETIS_V3_PartGeomKway (gkmetis.c:90)
17223	17223	libparmetis_Coordinate_Partition (xyzpart.c:58)
17223	17223	libparmetis_PseudoSampleSort (xyzpart.c:556)

Evaluate

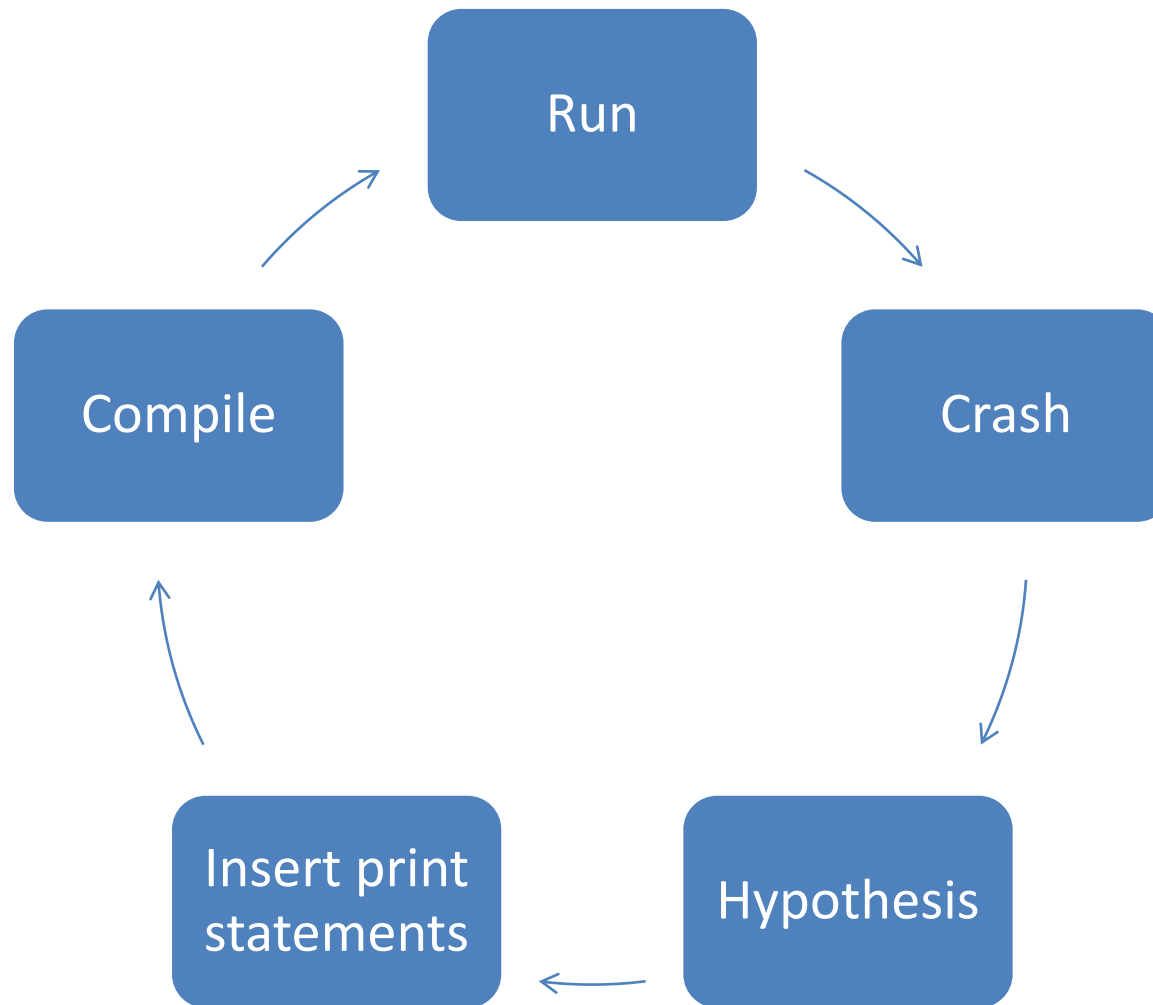
Expression	Value
i * ntsamples	-212322546

Type: int
Range: from -2147259746 to -12282046
49/17223 processes equal

7353 processes playing

Computer titan-ext7 Allinea DDT 4.2.1-36484 Sun Aug 10, 7:50 PM

Debugging in practice...



Some types of bug

Bohrbug	Steady, dependable bug
---------	------------------------


Heisenbug	Vanishes when you try to debug (observe)
-----------	------------------------------------------

Mandelbug	Complexity and obscurity of the cause is so great that it appears chaotic
-----------	---------------------------------------------------------------------------

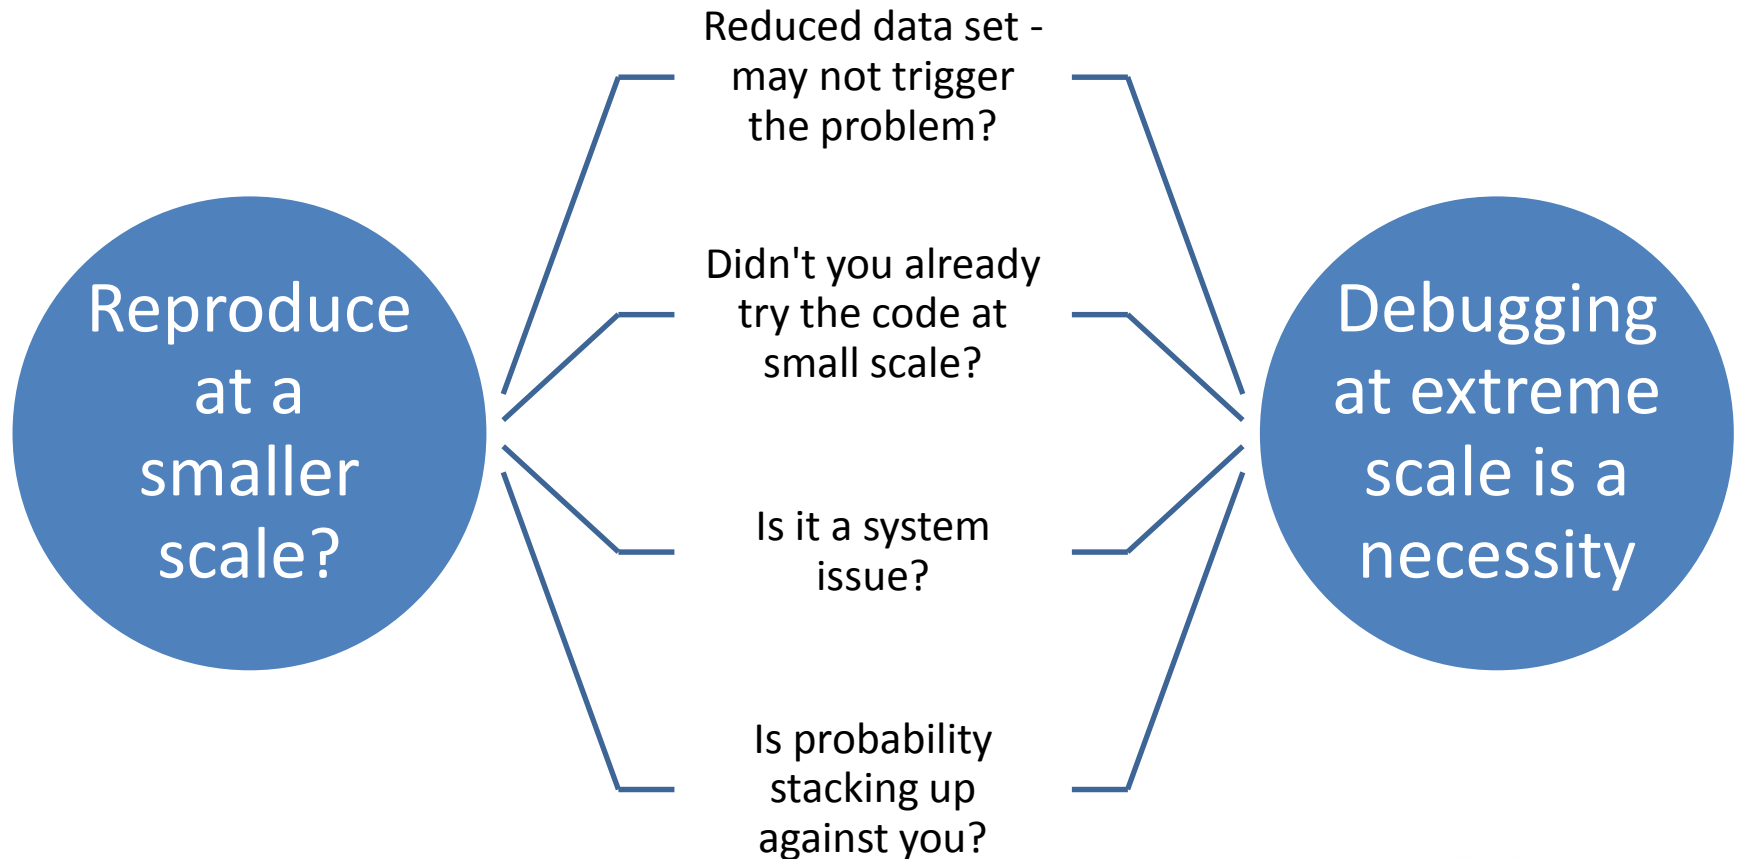
Schroedinbug	First occurs after someone reads the source file and deduces that it never worked, after which the program ceases to work
--------------	---------------------------------------------------------------------------------------------------------------------------

Print statement debugging


- The first debugger: print statements
 - Each process prints a message or value at defined locations
 - Diagnose the problem from evidence and intuition
- A long slow process
 - Analogous to bisection root finding
- Broken at modest scale
 - Too much output – too many log files



Bug fixing as scale increases



Extreme machines are everywhere



Machine
sizes are
exploding

Software
scale grows
as machines
grow

Titan and Mira

Titan

- 18,688 nodes
- 18,688 NVIDIA Kepler K20 GPUs
- 299,008 CPU cores
- 50,233,344 CUDA cores

Mira

- 49,152 nodes
- 786,432 cores
- 3,145,728 hardware threads

Does the printf workflow “work”?

ALCF, OLCF and Allinea deliver



2009 - Allinea and Oak Ridge begin collaboration to provide super-Petascale debugging




2010 - Allinea and Argonne collaboration to extend scaling to BlueGene systems




2013 - Mira and Titan full size debugging in place

What you should expect (demand!) for debugging at scale




Scalability

- A debugger that works to at least as high a scale as you need



Hardware and software support

- Whatever software you use and wherever you use it – the debugger supports it



Assistance

- Debugger is installed, configured, and documented – with site experts and training

Allinea DDT

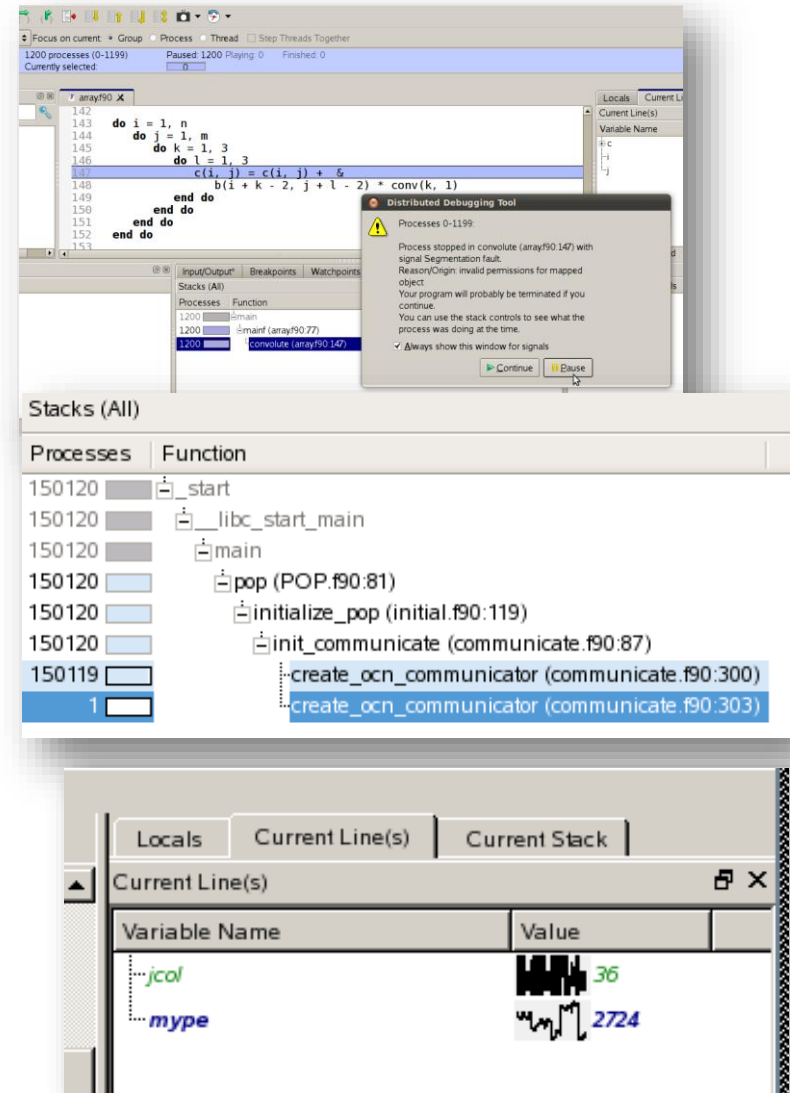
Fix software problems, fast

- **Powerful graphical debugger designed for :**
 - C/C++, Fortran, UPC, ...
 - MPI, OpenMP and mixed-mode code
 - Accelerators and coprocessors: CUDA and Intel Xeon Phi
- **Unified interface with Allinea MAP :**
 - One interface eliminates learning curve
 - Spend more time on your results
- **Slash your time to debug**
 - Reproduces and triggers your bugs instantly
 - Helps you easily understand where issues come from quickly
 - Helps you to fix them as swiftly as possible



Allinea DDT: Scalable debugging by design

- **Where did it happen?**
 - Allinea DDT leaps to source automatically
 - Merges stacks from processes and threads
- **How did it happen?**
 - Some faults evident instantly from source
- **Why did it happen?**
 - Real-time data comparison and consolidation
 - Unique “Smart Highlighting” – colouring differences and changes
 - Sparklines comparing data across processes
- **Force crashes to happen?**
 - Memory debugging makes many random bugs appear every time



The scalable print alternative

The scalable print alternative

Stop on variable change


Stop on variable change

Static analysis warnings on code errors

Static analysis warnings on code errors

Detect read/write beyond array bounds


Detect read/write beyond array bounds



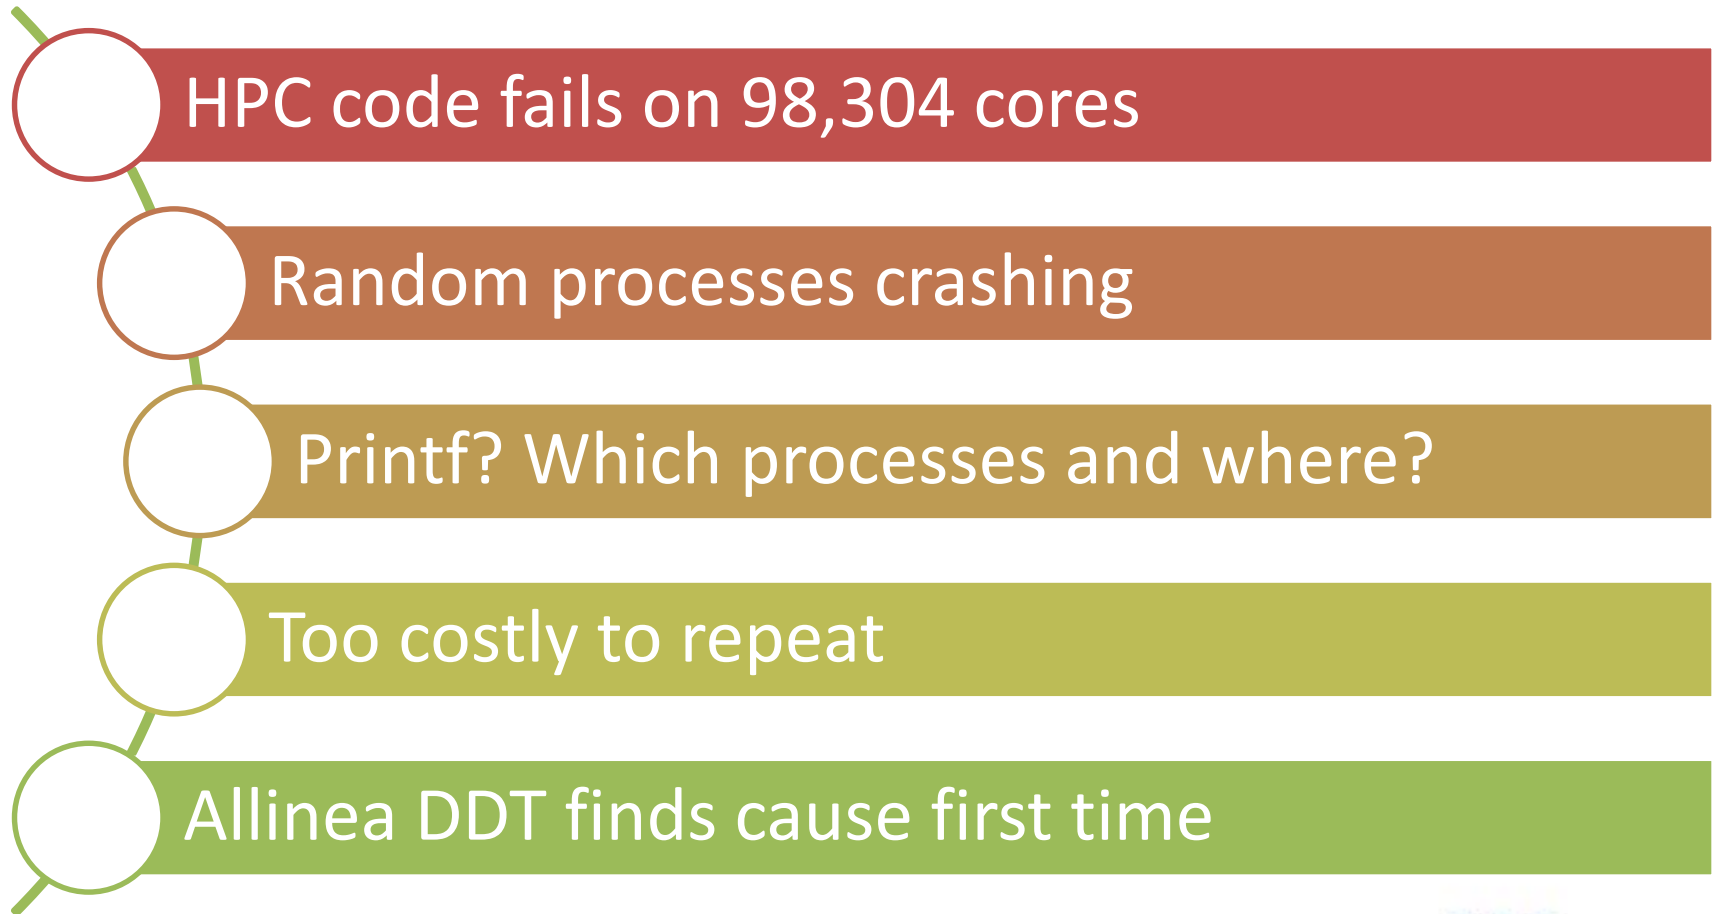
Detect stale memory allocations

Beneath the Petascale Allinea DDT

- Scalable tree network
 - Sends bulk commands and merge responses
 - Aggregations maintain the essence of the information
 - Step 100,000 processes?
100-150ms
- Usability matters
 - The interface is as important as the speed
 - Focus on scalable components




Example – ORNL's Titan



Offline debugging

- Interactive access difficult
- Use offline mode
 - Submit and forget
- Post-mortem analysis


► Stack for process 0

9  01:42.176 1-15 Process stopped in sched_yield (syscall-template.S:82) with signal SIGSEGV (Segmentation fault). Reason/Origin: kill, sigsend or raise
Your program will probably be terminated if you continue.
You can use the stack controls to see what the process was doing at the time.

▼ Stacks

Processes	Threads	Function
1-15	15	aplu (lu.f:118)
1-15	15	ssor (ssor.f:131)
1-15	15	blts (blts.f:55)
1-3,5-7,9-11,13,15	11	exchange_1 (exchange_1.f:37)
1-3,5-7,9-11,13,15	11	pmpl_recv
1-3,5-7,9-11,13,15	11	PMPI_recv
1-3,5-7,9-11,13,15	11	nca_pal_obl_recv
1-3,5-7,9-11,13,15	11	opal_progress
1-3,5-7,9-11,13,15	11	sched_yield (syscall-template.S:82)
4,8,12,14	4	exchange_1 (exchange_1.f:55)
4,8,12,14	4	pmpl_recv
4,8,12,14	4	PMPI_recv
4,8,12,14	4	nca_pal_obl_recv
4,8,12,14	4	opal_progress
4,8,12,14	4	sched_yield (syscall-template.S:82)

► Stack for process 1


10  01:46.488 n/a Every process in your program has terminated.

Messages Tracepoints Output

Tracepoints

#	Time	Tracepoint	Processes	Values
1	00:13.451	subdomain.f:96	0	jend: — 0 ny: — 9
2	00:13.453	ssor.f:177	0-15	delunm(5): — 0.25
3	00:13.455	ssor.f:177	0-15	delunm(5): — 0.25

Example – ANL Mira



Interlude: Local Demonstration

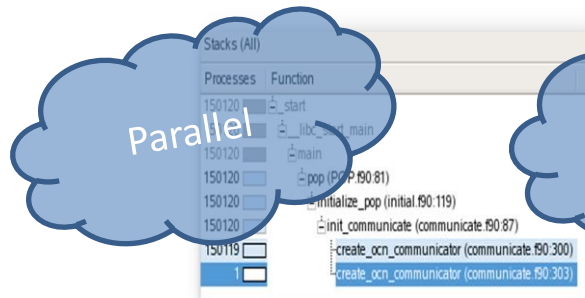
- Simple persistent hanging
 - Stepping through a code
- Process count dependent hanging:
 - Attaching to the running job

Getting started on Titan

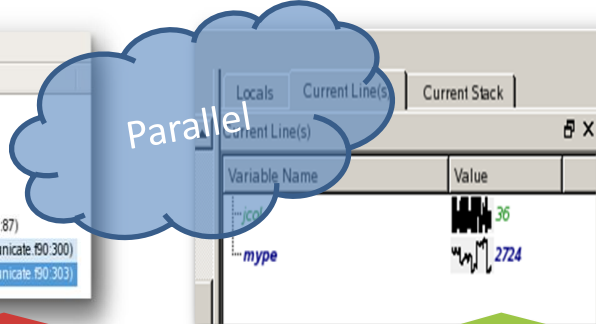
- How?
module load ddt
ddt
- Congratulations, you are now ready to debug.

Getting started on Mira/Tukey

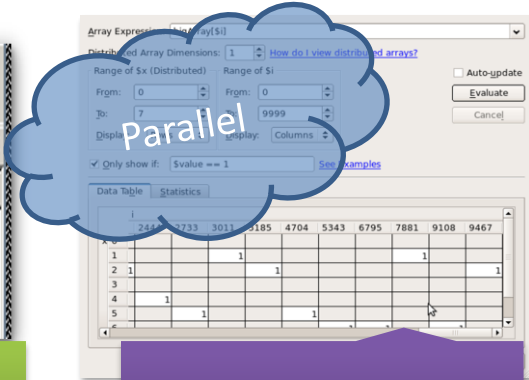
- Install local client on your laptop
 - www.allinea.com/downloads
 - Linux – installs full set of tools
 - Windows, Mac – just a remote client to the remote system
 - Run the installation and software
 - “Connect to remote host”
 - Hostname:
 - username@cetus.alcf.anl.gov
 - username@tukey.alcf.anl.gov
 - Remote installation directory: /soft/debuggers/ddt
 - Click Test
- Congratulations you are now ready to debug.



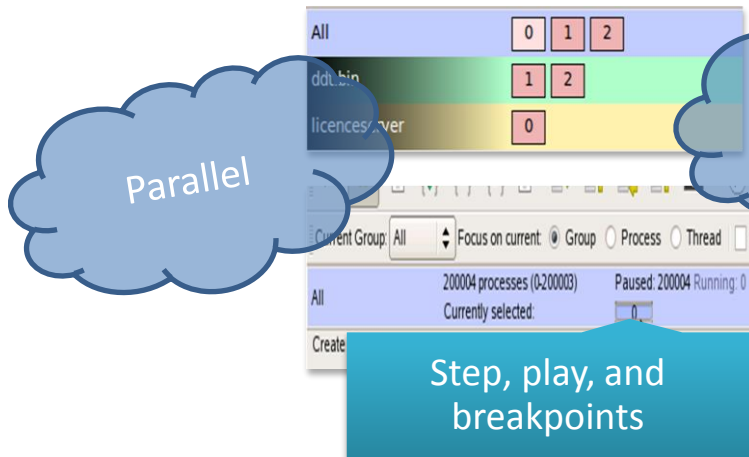
Parallel stack view



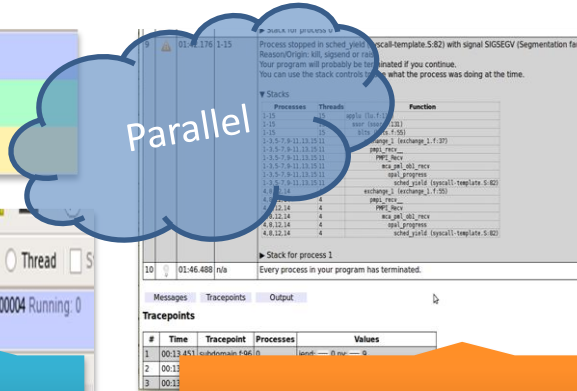
Automated data comparison: sparklines



Parallel array searching



Step, play, and breakpoints



Offline debugging

Summary



Debugging at scale is not difficult

- 300,000 cores is as easy as 30 cores
- The user interface is vital to success

Debugging at scale is not slow

- High performance debugging – at Mira and Titan scale
- Logarithmic performance

Stable, in production and well supported

- Routinely used over 100,000 cores